

SANDIA REPORT

SAND2017-11030

Unlimited Release

Printed October 2017

FY17 Status Report on the Computing Systems for the Yucca Mountain Project TSPA-LA Models

Teklu Hadgu, Gordon Appel, John Reynolds and Jason Garland

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525

Approved for public release; further dissemination unlimited.



Sandia National Laboratories



Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



FY 17 Status Report on the Computing Systems for the Yucca Mountain Project TSPA-LA Models

Teklu Hadgu, Gordon Appel, John Reynolds and Jason Garland

Nuclear Waste Disposal research and Analysis
Sandia National Laboratories
MS 0747
P.O. Box 5800
Albuquerque, New Mexico 87185

Abstract

Sandia National Laboratories (SNL) continued evaluation of total system performance assessment (TSPA) computing systems for the previously considered Yucca Mountain Project (YMP). This was done to maintain the operational readiness of the computing infrastructure (computer hardware and software) and knowledge capability for total system performance assessment (TSPA) type analysis, as directed by the National Nuclear Security Administration (NNSA), DOE 2010. This work is a continuation of the ongoing readiness evaluation reported in Lee and Hadgu (2014), Hadgu et al. (2015) and Hadgu and Appel (2016). The TSPA computing hardware (CL2014) and storage system described in Hadgu et al. (2015) were used for the current analysis. One floating license of GoldSim with Versions 9.60.300, 10.5, 11.1 and 12.0 was installed on the cluster head node, and its distributed processing capability was mapped on the cluster processors. Other supporting software were tested and installed to support the TSPA-type analysis on the server cluster. The current tasks included preliminary upgrade of the TSPA-LA from Version 9.60.300 to the latest version 12.0 and address DLL-related issues observed in the FY16 work. The model upgrade task successfully converted the Nominal Modeling case to GoldSim Versions 11.1/12. Conversions of the rest of the TSPA models were also attempted but program and operational difficulties precluded this. Upgrade of the remaining of the modeling cases and distributed processing tasks is expected to continue. The 2014 server cluster and supporting software systems are fully operational to support TSPA-LA type analysis.

ACKNOWLEDGMENTS

The authors would like to express their gratitude to Tito Bonano (6220) and Kevin McMahon (6222) for their technical interest and programmatic support.

Contents

FY 17 Status Report on the Computing Systems for the Yucca Mountain Project TSPA-LA Models	3
Acknowledgments	4
Tables.....	5
Nomenclature.....	6
1. Introduction	7
2. the tspa computing system	9
2.1. The TSPA Server Cluster Hardware (CL2014)	9
2.2. TSPA Server Cluster (cl2014) Hardware Issues.....	9
3. TSPA-LA Model Upgrade	9
3.1. TSPA-LA Model Conversion	9
3.2. Distributed Process	11
4. Conclusion and Future Work	12
4.1. Computing System.....	12
4.2. Execution of TSPA-LA model on the CL2014 Server Cluster.....	13
5. References	14
Appendix A. Summary Report of Hardware Malfunction and Remediation of the Problem	15
A1. TSPA Cluster (cl2014) Apparent Power and Thermal Issues in the Spring of 2017	15
A1.1 Background.....	15
A1.2 Power and Thermal Issues Timeline.....	16
A1.3 Discussion.....	17
Appendix B: Summary Report on Chronology of Troubleshooting Hardware Issues by Angela Calabeza.....	18
Appendix C: Summary Report of FY17 Activities from John Reynolds.....	20
Distribution.....	25

TABLES

Table A1.1 - Hardware and Software Evolution of TSPA Cluster.....	15
Table C 1 - Summary of FY17 Activities by John Reynolds	20

NOMENCLATURE

DLL	Dynamically Linked Libraries
DOE	Department of Energy
DP	Distributed Processing
DTN	Data Tracking Number
GB	Gigabyte
GTG	GoldSim Technology Group
LA	License Application
RAM	Random Access Memory
SNL	Sandia National Laboratories
TSPA	Total System Performance Assessment
YMP	Yucca Mountain Project

1. INTRODUCTION

Sandia National Laboratories (SNL) continued evaluation of the total system performance assessment (TSPA) computing systems for the previously considered Yucca Mountain Project (YMP). This was done to maintain the operational readiness of the computing infrastructure (computer hardware and software) and knowledge capability for total system performance assessment (TSPA) type analysis, as directed by the National Nuclear Security Administration (NNSA), DOE 2010. This work is a continuation of the ongoing readiness evaluation reported in Lee and Hadgu (2014), Hadgu et al. (2015) and Hadgu and Appel (2016). The TSPA computing hardware (CL2014) and storage system described in Hadgu et al. (2015) were used for the current analysis. One floating license of GoldSim with Versions 9.60.300, 10.5, 11.1 and 12.0 was installed on the cluster head node, and its distributed processing (DP) capability was mapped on the cluster processors. Other supporting software were tested and installed to support TSPA-type analysis on the server cluster. The current tasks included maintenance of the TSPA-LA system (hardware and software), and preliminary upgrade of the TSPA-LA from Version 9.60.300 to the latest version 12.0. The cluster cl2014 is currently running smoothly after having problems with operation of part of the compute nodes. Versions 9.60.300, 10.5, 11.1 and 12.0 are currently operational on the cluster. The 2014 server cluster and supporting software systems are now fully operational to support TSPA-LA type analysis with the GoldSim version used for the license application.

The model upgrade task continued from the work conducted in FY16. In FY17 the latest version of GoldSim, Version 12.0, (GoldSim, 2017) was installed on the cluster. To date the Nominal Modeling case was successfully converted to GoldSim Versions 10.5, 11.1 and 12.0. Upgrade of the remaining of the modeling cases and distributed processing tasks will continue.

In FY15, FY16 and FY17 SNL conducted evaluation of the TSPA computing system to verify the readiness of the capability to perform TSPA-type analysis of the Yucca Mountain repository following the 2014 server replacement. The reports by Hadgu et al. (2015) and Hadgu and Appel (2016) documented the work performed to achieve and maintain the readiness of the computing infrastructure (computer hardware and software) and knowledge capability to perform TSPA-type analyses. The reports provided details of specifications of the 2014 computer hardware, the evaluation of the required components of the hardware and software systems, as well as the instructions to setup and conduct the TSPA-LA type simulations and post-processing of the model output. This report is a continuation of the work performed in 2015 and 2016.

As was done in previous work (Lee and Hadgu, 2014, Hadgu et al., 2015, Hadgu and Appel, 2016) one of the goals of this work is to demonstrate the readiness of the 2014 hardware and software systems. This is to insure that the computing system can support reliable execution of the TSPA-LA models and post-processing of the model output. This includes completing tasks started in FY16. The other goal of this work is to start upgrading the TSPA-LA from the original GoldSim version 9.60.300 to the latest version (12.0). The following main topics were identified for the current investigation to evaluate the status of the TSPA-LA model capability.

- Maintaining readiness of the TSPA-LA
- To begin upgrading of the TSPA-LA models with conversion from GoldSim 9.60.300 to GoldSim version 12.0 (GoldSim, 2007, 2017).

2. THE TSPA COMPUTING SYSTEM

2.1. The TSPA Server Cluster Hardware (CL2014)

The TSPA computing system, which constitutes the hardware and software, is discussed in great detail in Hadgu et al. (2015). The new TSPA cluster (CL2014) consists of a total of 32 Dell PowerEdge R620 servers, each with 3.0 GHz Intel® Xeon® E5-2690 dual quad-core processors (20 processors per server) and 128 GB RAM. Thus, the TSPA server cluster has a total of 640 processors.

The 2014 servers reside on the Sandia DMZ domain and are running under the Windows Server 2012 r2, 64-bit operating system. The system was optimized for installation and execution of the GoldSim software required to run the GoldSim distributed processing module utility (GoldSim 2010, 2017). The distributed processing module utility is a program extension to GoldSim which allows use of multiple computers connected over a network to share the computational burden of a Monte Carlo simulation. The module is the essential feature to efficiently manage and execute multiple realizations of the TSPA-LA model run on the cluster processors.

In the 2014 configuration of the total of 32 blade servers, one blade server is used as the head node, and 31 servers are used as the compute nodes dedicated to run GoldSim-based TSPA models.

2.2. TSPA Server Cluster (cl2014) Hardware Issues

The TSPA cluster (cl2014) had hardware issues for parts of FY17. During model runs several compute nodes at a time were not functional. The hardware problems are discussed in detail in Appendix A. These problems persisted for a considerable amount of time, and thus affected the progress of model upgrade. The cluster is currently functioning properly.

3. TSPA-LA MODEL UPGRADE

3.1. TSPA-LA Model Conversion

The 2008 TSPA-LA simulations that were an integral part of the license application were based on GoldSim 9.60.300. Since then several upgrades have been made to GoldSim software. Because of the possibility that the older version will not be supported in the near future by GoldSim as well as the WINDOWS operating system, we have started upgrading the TSPA-LA to later GoldSim versions. Upgrading was started in FY16 (Hadgu and Appel, 2016). The model upgrade work has continued in FY17. The following is a summary of the status of the work. Part of the upgrade work reported in the FY16 report is summarized here for completeness.

As documented in the FY16 report (Hadgu and Appel, 2016), the Nominal Modeling case was upgraded in two stages: GoldSim 9.60.300 to 10.5 and then 10.5 to 11.1. This was done to simplify and manage the conversion process. Opening of the 2008 TSPA-LA Nominal Modeling case in GoldSim 11.1 resulted in a number of error messages. A list of the error messages and the solutions applied to mitigate the errors, with the exception of errors related to the distributed process system, are described in detail in the FY16 report and summarized below. Solutions to problems with the distributed processing system are described in this report in Section 3.2.

The conversion process of the Nominal Modeling case from GoldSim 9.60.300 to 11.1 required addressing the following issues:

- 1) SubModel output interface conversions: Two SubModels were flagged for conversions made to the outputs on their output interface. The conversion messages note that this version of GoldSim “*changes how simulation data are exported from SubModel elements*”. The SubModels are Aleatory_Params (located in \Time_Zero) and Epistemic_Params (located in \Epistemic_Uncertainty). The conversion messages note that all outputs were converted successfully. This conversion is not expected to affect model results.
- 2) Table log interpolation changed to linear interpolation: Five Lookup Table elements were flagged because they “*used to do log interpolation on the independent variable*”. As of GoldSim version 11.0, the log interpolation option is no longer available for Lookup Tables. The conversion message noted that these were converted to linear interpolation.
- 3) Conversion of custom unit built-in SI unit: One conversion message noted the “*unit cdeg is not defined*”. The cdeg is a custom unit defined in the TSPA model. The conversion message presents the following question and provides ‘Yes’ and ‘No’ options: “*Did you mean the SI unit Cdeg?*” The ‘Yes’ option was selected.
- 4) Time Series Definition: The external exchange format for the Time Series data type has changed. Thus, the old DLL (TS_Proc.dll) function does not operate correctly. This DLL has now been replaced with the latest DLL (TSProc.dll).
- 5) Custom Resampling Logic: The resampling logic for a Stochastic element correlated to another element has been changed. In the current version these options are mutually exclusive, and thus the correlation option has been changed to “not correlated”. This affected eight Stochastic elements.
- 6) Recording Time Series Workaround: Located in the Submodel ‘EBS_Submodel’ (at \Global_Inputs_and_Calcs\Global_Events\Seismic_Scenario\Model_Input_Seismic\Model_Feeds_Seismic\Aleatory_Feeds_Seismic), there is a Time Series element, ‘Seismic_Event_Occurs’, that generates an error when the TSPA model is run in GoldSim version 11.1.5. The error message is “No data in Time Series ...”. This issue was addressed by implementing a workaround described in detail in Appendix B.
- 7) Running with Distributive Process: Running the Nominal Modeling case on many processors results in error messages. These messages seem to be related to unavailability of certain input files when they are needed by a process.

In FY16 all of the error messages related to the Nominal Modeling case, except for the distributed process, were addressed. Changes were made to the modeling case to address the errors. Some of these changes have the potential to affect output results. A study of the effect of the changes on output is on-going. The following changes resulting from the conversion process could affect model results:

- 1) The correlation option for eight Stochastic elements was changed to ‘not correlated’.
- 2) The log interpolation setting for five Lookup Tables was converted to linear interpolation.
- 3) A custom unit (cdeg) was converted to the built-in SI unit Cdeg.

The following changes from the conversion process should not affect model results:

- 1) A change in the external exchange format for the Time Series data type required the use of a different DLL (TSProc.dll) in place of TS_Proc.dll.
- 2) Outputs on the output interface of two Submodel elements were converted.

Also, the workaround implemented in the 11.1.5 version of the model is not expected to impact model results (see Hadgu and Appel, 2016).

As shown in the FY16 report the GoldSim 11.1 version of the Nominal Modeling case was successfully run to completion. Comparison of total dose output showed minor differences mainly at early time, at very low total dose. Note that some of the changes made to the Nominal Modeling case, such as changing the resampling plus correlation option, did not affect the results because they were not needed for this modeling case. In FY17 the Nominal Modeling case was successfully upgraded to GoldSim 12.0 with no additional error messages. Conversion of the rest of the TSPA-LA models is in progress. The conversion effort for the rest of the modeling cases is done in two stages as was done for the Nominal Modeling case: GoldSim 9.60.300 to 10.5 and 10.5 to 12.0. First each of the modeling cases was opened in GoldSim 10.5 to record any errors encountered. Below is a summary of conversion errors encountered for each modeling case while converting to GoldSim 10.5.

Drip Shield Early Failure Modeling Case (3,000 realizations) - Error in element
'Mean_CDSP_Thk_Out': "Table data has no rows defined"

Waste Package Early Failure Modeling Case (6,000 realizations) - Error in element
'Mean_CDSP_Thk_Out': "Table data has no rows defined"

Seismic Ground Motion Modeling Case (9,000 realizations) - Error in element
'Solubility_U_CSNF': value out of range. Similar error with other elements.

Seismic Fault Displacement Modeling Case (10,800 realizations) - Error in element
'CSNF_U_Schoepite_TypeIII_Cap': value out of range. Similar error with other elements.

Igneous Intrusion Modeling Case (3,000 realizations) - Error in element
'Mean_CDSP_Thk_Out': "Table data has no rows defined"

Volcanic Eruption Modeling Case (12,000 realizations) - No conversion errors.

Human Intrusion Scenario (9,000 realizations) - Error in element 'Mean_CDSP_Thk_Out':
"Table data has no rows defined"

It appears that the above errors are the result of the conversion changes such as removing the resampling plus correlation option. Efforts were made to identify the impact of the software changes by making additional changes in the models and also by modifying DLLs. Activities related to these efforts are outlined in Appendix B. Further work will be needed to solve DLL issues and to find alternative methods of achieving functions implemented in GoldSim 9.60.300, but no longer available in later versions. The conversion effort has continued with support from the GoldSim Technology Group (GTG).

3.2. Distributed Process

All TSPA modeling cases are run under GoldSim distributed processing system, where simulation is done on a user specified number of processors. As specified in the FY16 report (Hadgu and Appel, 2016) problems were encountered with the distributed processing system when GoldSim 11.1 was used. Details of the problems are documented in Hadgu and Appel (2016, Appendix C). The problems seem to be related to file access issues. One of the problems is that when running the TSPA models one realization could be assigned to more than one compute nodes causing file access errors. This implementation was added to GoldSim in support of the Yucca Mountain Project. The latest version of GoldSim (GoldSim 12.0) eliminates the implementation. The distributed processing system in GoldSim 12.0 has now been tested. The TSPA Nominal Modeling case was run with 300 processors. The file access errors encountered in GoldSim 11.1 related to the distributed processing system are now absent. This addresses the malfunction. Any adverse effects of removing the implementation will be monitored.

As reported in Hadgu et. al. (2015) one other issue related to the GoldSim distributed processing system was the maximum number of processors allowed per machine. In GoldSim 9.60.300 a maximum of 10 processors per machine was allowed. This limited the number of realizations that could be run on Cluster cl2014 to a maximum of 310 (i.e. 31 compute nodes x 10). This requirement has been relaxed in later versions of GoldSim. Since the cluster has 20 processors per machine, a total of up to 620 processors are available. To investigate any possible improvements, the distributed processing system of GoldSim 12.0 was tested using the Nominal Modeling case with 20 processors per machine. The run was not able to complete. Various numbers of processors per machine were then tested to identify the maximum possible. The testing results showed that the maximum number of processors per machine that could be achieved was 11. Further testing will be required to identify the cause of the limitations.

4. CONCLUSION AND FUTURE WORK

The purpose of this work is to evaluate and maintain operational readiness of the computing infrastructure (computer hardware and software) and knowledge capability to perform TSPA-LA type analyses. To achieve the objective a multi-year on going work is in progress. Previous work has been documented in Lee and Hadgu (2014), Hadgu et al. (2015) and Hadgu and Appel (2016). The following tasks were conducted in FY17.

- Evaluation and maintenance of the CL2014 TSPA server cluster system to support reliable executions of the TSPA-LA models and associated analysis and calculations. Hardware malfunctions have been addressed and the cluster is now fully operational.
- Retrieval of the TSPA-LA model files required input files and other associated files of the TSPA-LA modeling cases.
- Execution of the TSPA-LA model on the TSPA cluster servers (CL2014), ensuring reliable run executions utilizing the GoldSim distributed processing module and reproducible stochastic sampling schemes (GoldSim 2007). Several of the TSPA models were re-run in FY17 using GoldSim 9.60.300 to test the TSPA cluster. The execution times matched previous runs.

Future work will include upgrading of the TSPA-LA models to GoldSim 12.0 (the latest version) and conducting the tasks that were executed in FY15 and FY16.

4.1. Computing System

The TSPA cluster (CL2014) consists of a total of 32 Dell PowerEdge R620 servers, each with 3.0 GHz Intel® Xeon® E5-2690 dual quad-core processors (20 processors per server) and 128 GB RAM. Thus, the TSPA server cluster has a total of 640 processors. The cluster runs under Windows Server 2012 r2, 64-bit operating system. The operating system was optimized for installation and execution of the GoldSim software required to run the GoldSim distributed processing module utility (GoldSim 2010).

The TSPA-LA models (SNL 2008) were developed with GoldSim Version 9.60.300. A floating license of Version 9.60.300 (SP3) has been installed on the cluster head node, and its distributed processing capability was mapped on the cluster processors. GoldSim Versions 10.5, 11.1, and 12.0 are also installed on the cluster as part of TSPA-LA model upgrade to the latest version of GoldSim.

As described in Section 2.1 and detailed in Appendix A, the TSPA cluster had hardware problems for part of the year. The malfunction has now been addressed and the cluster is currently fully operational. Addressing the malfunction slowed the model upgrade work.

4.2. Execution of TSPA-LA model on the CL2014 Server Cluster

As documented in Hadgu et al. (2015) and Hadgu and Appel (2016) the latest TSPA-LA models for individual modeling cases retrieved from DTN MO0710ADTSPAWO.000 (GW Modeling cases (v5.005) without Final Documentation) were executed on the CL2014 TSPA server cluster to evaluate performance of the cluster using GoldSim 9.60.300. All runs were executed on multiple processors on the cluster servers utilizing the GoldSim distributed processing modules (GoldSim 20010, 2017), and all runs were completed successfully.

As the GoldSim model has been progressively upgraded from previous versions one of the tasks identified was to start converting the TSPA-LA models to later versions so that it is up-to-date. This work was started in FY16 as documented in Hadgu and Appel (2016). For FY16 the Nominal Modeling case was upgraded to GoldSim version 10.5 and 11.1. The 2016 work identified issues with conversion including use of the distributed process system. Upgrading of the TSPA-LA to later versions continued in 2017. To date the latest version of GoldSim (12.0) has been installed on the cluster and the Nominal Modeling case has been successfully run using the new version. Issues with the distributed processing system identified in GoldSim 11.1 has been addressed in GoldSim 12.0. Upgrading of the rest of the modeling cases to GoldSim 10.5 and then 12.0 has started and is discussed in Section 3.1. The task was delayed due to cluster hardware malfunction. Upgrading of the modeling cases resulted in error messages related to changes in the later versions of GoldSim. Some of the errors seem to be related to unavailability of functions that were part of GoldSim 9.60.300, the version that was used for the TSPA-LA. Work is in progress to address these issues and DLL-related issues with support of GoldSim Technology Group.

5. REFERENCES

GoldSim (GoldSim Technology Group) 2007. User's Guide, GoldSim Probabilistic Simulation Environment. Version 9.60. Two volumes.; User's Guide, GoldSim Contaminant Transport Module. Version 4.20; User's Guide, GoldSim Distributed Processing Module, Version 9.60. Issaquah, Washington: GoldSim Technology Group, March 2007.

GoldSim 2010. User's Guide, GoldSim Probabilistic Simulation Environment. Version 10.50, Two volumes; User's Guide, GoldSim Contaminant Transport Module. Version 6.0; User's Guide, GoldSim Distributed Processing Module, Version 10.50. Issaquah, Washington: GoldSim Technology Group, December 2010.

GoldSim 2017. User's Guide, GoldSim Probabilistic Simulation Environment. Version 12.0, Two volumes; User's Guide, GoldSim Contaminant Transport Module. Version 12.0; User's Guide, GoldSim Distributed Processing Module, Version 12.0, 22500 SE 64th Place, Suite 240, Issaquah, Washington: GoldSim Technology Group, February 2017.

Hadgu, T., Appel, G., Malashev, A., and Payne, C. 2015. Reproduction of the Yucca Mountain Project TSPA-LA model Runs using TSPA Computing Systems, SAND2015-7834.

Hadgu, T., and Appel, G. 2016. Reproduction of the Yucca Mountain Project TSPA-LA Uncertainty and Sensitivity Analyses and Preliminary Upgrade of Models, SAND2016-9957

Lee, J. and Hadgu, T. 2014. Evaluation of the Computing Systems for Yucca Mountain Repository TSPA-LA Model Operational Readiness, Sandia National Laboratories, Albuquerque, New Mexico. SAND2014-18178.

Appendix A. Summary Report of Hardware Malfunction and Remediation of the Problem

A1. TSPA Cluster (cl2014) Apparent Power and Thermal Issues in the Spring of 2017

A1.1 Background

In 2010 the YM TSPA-LA hardware and software was moved from the YMP in Las Vegas, NV to Sandia National Laboratories (SNL) in Albuquerque, NM. SNL has maintained the operational readiness of the computing infrastructure (computer hardware and software) and knowledge capability for total system performance assessment in the license application (TSPA-LA), in the event of future needs related to the YM Project. Table 1 summarizes the history of hardware and software used for the YM TSPA-LA.

Table A1.1 - Hardware and Software Evolution of TSPA Cluster

Cluster Vintage	GoldSim Version	Initial Use	Current Use	Projected Use
2006	9.60.300	TSPA-LA @ YMP	Repurposed	Retire
2010 (hn01snlntz)	9.60.300	TSPA-LA @ SNL	Repurposed	None
2015 (cl2014)	9.60.300	TSPA-LA @ SNL	TSPA-LA @ SNL	Active
2015 (cl2014)	11.1 / 12	TSPA-LA @ SNL	Not Applicable	Not Applicable

In early FY15 the then existing TSPA-LA computational cluster (cl2010) was replaced with the cl2014 hardware. TSPA-LA modeling cases were tested and verified for model reproducibility on the cl2014 server. As before, all test runs were executed on multiple processors on the server cluster utilizing the GoldSim distributed processing capability, and all runs were completed successfully. The FY 15 and FY16 work is documented in two SAND reports^{1,2}.

The FY17 scope of work was intended to consist of:

¹ Hadgu, T., Appel, G., Malashev, A., and Payne, C. 2015. Reproduction of the Yucca Mountain Project TSPA-LA model Runs using TSPA Computing Systems, SAND2015-7834.

² Hadgu, T., and Appel, G 2016. Reproduction of the Yucca Mountain Project TSPA-LA Uncertainty and Sensitivity Analyses and Preliminary Upgrade of Models, SAND2016-9957

- Continue evaluation of the extent of DLL problems needing resolution to run the TSPA-LA on GoldSim Version 11.1/12.
- Resolve the distributed processing problems to the extent practicable and if possible confirm that TSPA-LA runs on GoldSim Version 11.1/12
- Begin verification of the TSPA-LA results achieved running GoldSim Version 11.1/12 on cl2014 hardware, by comparison to the results used to support the YM License Application.
- Document the comparison of the results achieved running GoldSim Version 11.1/12 on cl2014 to the results used to support the YM License Application.
- Perform and verify TSPA-LA Uncertainty and Sensitivity analysis GoldSim Version 11.1/12
- Perform Benchmark testing (computational efficiency) on GoldSim Version 11.1/12

A1.2 Power and Thermal Issues Timeline

After its installation in early FY15, cl2014 has functioned as expected throughout FY15, FY16, and until late April of FY17. Significant events related to this issue are provided below.

4/25/17 – TSPA Technical Lead (Teklu Hadgu) notices TSPA cluster (cl2014) having power off issues during TSPA runs

4/25-28/17 – Attempts by IT server technician (Jason Garland) to resolve issue not successful

4/28/17 – Teklu Hadgu submits ticket to ITSM Production (INC000002011791) “Power issues to CL2014”

5/3/17 – Teklu Hadgu inquires after resolution of ticket, informed he will be contacted and later advised to put another service request in to the annex team, DCFIT.

5/4/17 – IT server technician is assured electricians have checked the power and suggest it could be the cl2014 equipment

5/11-22/17 – Testing of power-related issues continue, as well as looking into replacement of parts under warranty from the supplier (Wildflower/Dell) of one or more power distribution units (PDUs).

5/18/17 – TSPA Technical Lead and Project Manager inform 8842 Manager (Kevin McMahon) of situation. TSPA Technical Lead scheduled to be out of office next 2 weeks (Conference and Vacation).

5/22/17 Mission Computing Services (MCS) Service Manager 9322 (Angela Calabaza) introduced into situation

5/30/17 – TSPA Project Manager (Gordon Appel) follows up with IT server technician in view of apparent lack of progress. Told by IT it was still working with DELL.

5/30-6/6/17 – TSPA Project Manager has several conversations with MCS Service Manager stressing importance of restoring cl2014 to full service. MCS Service Manager makes Data Center Operations Manager 10774 (Tony Perea) aware of circumstances.

6/1/17 – TSPA Project Manager apprises 8840 Manager (Tito Bonano) of ongoing situation.

6/6/2017 – TSPA and IT meeting about needing a clearly thought out and methodical approach to identify and solve the problem.

6/7/17 – TSPA Technical Lead submits Nominal Case (300 realizations) to run on cl2014 on 50 processors. **Computation is successful.**

6/7/17 – TSPA Technical Lead submits Nominal Case (300 realizations) to run on cl2014 on 300 processors. **Computation is successful.**

6/7/17 – MCS Service Manager provides approach to TSPA Team that focuses on further investigation of power and thermal issues.

6/8/17 – TSPA Technical Lead submits Seismic Ground Motion Case (9000 realizations) to run on cl2014 on 300 processors. **Computation is not successful.**

A.1.3 Discussion

Presently, we are still experiencing power/thermal issues on the cl2014 server. This does not remove the ability to conduct the calculations; however, it does decrease the efficiency of doing so. The IT team has proposed an approach to further evaluate the issue though they are not certain about the cause of the problem and cannot assure its resolution.

Appendix B: Summary Report on Chronology of Troubleshooting Hardware Issues by Angela Calabeza

I would like to thank you for your patience while MCS troubleshooted your CL2014 cluster issues by working with our Sandia partners and vendors. Jason worked this issue for about 4 weeks with our Sandia partners and Dell support staff. Dell alleged power and/or thermal issues and our Sandia partners suspected hardware issues. After a delay in problem resolution, Angela worked with Tony Perea to get representatives from all concerned support teams together in a meeting (June 8th) to discuss resolution strategies.

We are unable to state with 100% certainty what caused the problems. We can speculate on likely causes and can list the corrective actions taken. Here are the list of things MCS and DCC/DCFIT did to ensure the cluster was working correctly and processes put in place to mitigate issues in the future.

Prior to the meeting DCFIT made the following updates:

- May 10th - Joey Sena informed Jason that they needed to replace the PDU going to the rack where the CL cluster resides. Once the PDU was replaced, DCFIT was going to monitor the PDU.
- May 16th - Dave Martinez stated the PDU in the rack was replaced. No power to the server was lost. The PDUs were balanced and there are no power issues.
 - The PDUs were redundant. Replacing one PDU, the cluster would have “failed over” to the second PDU and would have not lost any power. Depending on steps taken during the replacement action, it’s likely that either a brownout or surge condition occurred during the “fail over.” We think this may have caused some of the nodes to reset the BIOS and performance settings. Specifically, the processor fan speed was set to a low rpm configuration, which made thermal issues inevitable on the affected systems. Upgrading the BIOS, firmware, iDRAC, and IPERC likely resolved the fan speed issue.
 - June 29th – Jason sent a request to Dell to see if we could pull data as to when the BIOS and fan speed may have changed on the nodes.
 - Dell responded back and stated the reports were not that granular and we couldn’t see when BIOS or fan speed changed.
- June 6th – DCC and DCFIT pulled cables to setup the iDRAC configuration.

After the meeting, DCFIT took the following action items:

- Set up thermal readers
 - June 22nd – Thermal reading showed the temperature was fine on the front of the cluster. There were some warm spots on the back of the server but that’s when testing occurred and the server should warm up.
- Verify cooling and coils for the master and slave units to ensure there is steady temperature
 - June 22nd – Temperature was set and working correctly for the YMP room.

- Verify the air flow across and out of Annex.
 - June 22nd - Air flow was set and working correctly for the YMP room.
- During the past 3 months - DCFIT and DCC removed some servers from the YMP room.
 - The YMP room is running colder now and with better air pressure.

After the meeting, MCS performed the following upgrades to ensure the cluster was functioning:

- Work with Bob D'Spain on setting up iDRAC on CL2014.
 - June 21st - The iDRAC setup was completed. Jason was able to pull data from cluster. The data didn't show power or thermal issues.
- Verify the Watts on each node and send the information to the team.
 - June 21st - The wattage on the PSU have been verified to be 750 Watts. When I've talked to the Dell tech they said that any power issues would show up in the IDRAC reports I've given them. They now have a total of 13 reports with zero power issues detected by the IDRAC.
- Update the firmware, iDRAC, and iPERC on all nodes.
 - June 20th – Jason completed firmware updates on 30 nodes.
- Update the Bios on all nodes
 - June 20th – Jason completed BIOS updates on 14 nodes. These were the nodes failing.
- Verify the hyperthreading setup and fan configuration on all nodes
 - June 27 – fan configuration changes on 14 nodes. These were the nodes failing.
- Follow up with Dell to have them send reports of thermal failures. Once we have them they will be sent to the team.
 - June 7th - Jason provided DCFIT with the Dell logs stating thermal failures.
 - June 13th – DCFIT didn't see issues that would have been caused by Annex power or thermal issues.

Processes Put in Place:

- MCS will keep a change log of the CL cluster. This log will show any hardware failures, power outages, software updates, etc.
- MCS will perform configuration management of the server. This will document all troubleshooting steps and issues found during this exercise. This will be updated as changes and/or updates occur.
- MCS highly recommends installation of Nagios on the cluster. This free software enables real time, automated monitoring of environmental sensors and other cluster health indicators. Nagios will raise configurable alarms if thermal conditions leave nominal ranges. This will allow us to mitigate issues in the future before they become problems.

Thanks

Angela Calabaza, Service Manager
Mission Computing Services (9322)
Sandia National Laboratories
(505)284-2515

Appendix C: Summary Report of FY17 Activities from John Reynolds

Table C 1 - Summary of FY17 Activities by John Reynolds

Task Name	Description
Misc Support	TSPA troubleshooting
Miscellaneous overhead tasks	Discussed intern with Teklu Hadgu
	Discussed meeting scheduling with Teklu Hadgu
	Discussed TSPA schedule with Teklu Hadgu
	Impromptu TSPA model conversion meeting
	Located PSQ site for Teklu Hadgu. Also, reviewed his laptop choice upon request
	Quarterly 6222 SLA review
	Sat down with Kevin McMahon to discuss TSPA project scope change
	Sent Angela Calabaza notification of scope change on TSPA project
	Set up meeting with Jason P Garland, Teklu Hadgu, Goldsim staff
	Talked to Teklu Hadgu about WAPDEG and Jesse scheduling
	Troubleshoot cluster issues with GoldSIM
	Troubleshooting GoldSim launch failure issues
	Updated list of duties for 2018 SLA
	WAPDEG status meeting
SD - Other	Added additional error checking to WAPDEG DLL, sent to Teklu Hadgu for testing
	Answered Teklu Hadgu email about DLL with improved diagnostic output
	Built and sent corrected WAPDEG DLL to TH
	Built latest WAPDEG DLL
	Discuss GoldSim with Teklu Hadgu
	Discuss status with Teklu Hadgu
	Discuss test cases on multicore with Teklu Hadgu
	Discuss TSPA with Teklu Hadgu
	Discussed GoldSim issues with Teklu Hadgu
	First shot at getting WAPDEG to work with latest GoldSim version
	Grind on GoldSim issue, requested source code for multiple DLLs
	Met with Jesse Saunders about WAPDEG project
	Met with Teklu Hadgu to change GoldSim server behavior using import files

	Met with Teklu Hadgu about WAPDEG table not being written
	Responded to Goldsim support post by Ryan Roper
	Sent FileExists routine to Jesse B Saunders
	Sent Goldsim help center note indicating that broken correlations are the most serious matter facing us
	Sent issue report to GoldSim help center - Undefined fatal error running model
	Troubleshoot GoldSim error with Teklu Hadgu
	Troubleshoot GoldSim issue - TS_PROC_DLL function list updated, unspecified error
	Troubleshoot GoldSim models (uncertainty correlation issues)
	Troubleshoot issue with WAPDEG dll
	Troubleshoot issues found in first test of "new" WAPDEG DLL
	Troubleshoot TSPA issue with WAPDEG
	Troubleshoot TSPA model - interpolation on table lookups
	Troubleshoot TSPA model with Teklu Hadgu
	Troubleshooting GoldSim issue with Teklu Hadgu
	TSPA model conversion meeting
	TSPA Troubleshoot meeting
	WAPDEG - prepare for second try -- added sleeps
	WAPDEG troubleshooting
	Worked on GetThk_LA
	Worked on TSPA - troubleshoot reason WAPDEG could not write to files
	Worked on WAPDEG DLL
	Worked with Jesse on creating unique filename function
Software meeting/communication about software	Discuss GoldSim versions with Teklu Hadgu
	Discussed issues around Goldsim project with Teklu Hadgu- informal meeting
	Discussed schedule with Teklu Hadgu
	Goldsim meeting
	Impromptu meeting about TSPA model conversion
	Meeting with Goldsim company staff
	Met about YMP GS activities
	Met to discuss project and GoldSim 12

	Met with Teklu Hadgu and Jason Garland to discuss GoldSim server and licensing problems
	Met with Teklu Hadgu about new plan for fixing GoldSim issue
	Multiple preparatory meeting discussions and tasks
	Multiple troubleshooting meetings with Teklu Hadgu about GoldSim 12.0
	SLA meeting (YMP)
	TSPA Model Conversion meeting with Teklu Hadgu
System Administration Tasks, desktop	Helped out Teklu Hadgu with secure ssh solution
	Helped Teklu with installing new WinZip from SAMS
	Reviewed system quote for Teklu Hadgu
System Administration Tasks, server	8853 quarterly review meeting with Angela Calabaza
	Bi-weekly MCS IT Support team meeting - topics related to CL2014 system outage
	Cluster troubleshooting meeting with Angela Calabaza and various computer annex subject experts.
	CS2014 Cluster meeting - discuss official write-up of troubleshooting and corrective actions taken
	Discuss server issues with Teklu Hadgu
	Discussed CL2014 server issues with Angela Calabaza and Sheri Gibson
	Discussed CL2014 system issues with Teklu Hadgu and Jason Garland (separate times)
	Discussed MCS IT support team meeting with Teklu Hadgu
	Discussed server issues with Teklu Hadgu and Jason P Garland
	Discussed servers with Jason Garland
	Discussed test case submittal with Teklu hadgu
	Edited CL2014 troubleshooting and mitigation report
	Discussion with Teklu Hadgu about cause of thermal issues in CL2014 cluster.
	Launched test TSPA jobs on CL2014
	Meet about CL2014 server
	Meeting about cluster issues
	Met with Teklu Hadgu and Jason Garland about issues with the CL2014 cluster
	Miscellaneous troubleshooting of GoldSim runs

	Modified scripts for running beta GoldSim
	Requested Sysadmin declare network purpose for CL2014 systems
	Reviewed GoldSim issue discussion thread, determined that registry key needed to bypass redundant realization submission not visible to GoldSim user (thadgu) actually running .exe
	Short meeting with Teklu about CL2014 cluster
	Suggested escalation of issue with Teklu's cluster
	Troubleshoot GoldSim multicore submittal issues
	Troubleshoot issues on Teklu Hadgu's cluster
	Troubleshoot multiprocessing issue, showed Teklu Hadgu some things he can look at to assess performance
	Troubleshooting CL2014 cluster outage
	Troubleshooting meeting, Teklu Hadgu and Jason Garland
	Wrote email outlining plan of attack for getting Teklu's cluster back up
Technical Support	Discussed issues about Goldsim 12.x job submittal with Teklu Hadgu
	Helped Jesse with build troubleshooting
	Helped Teklu Hadgu get administrative access to install MatLab
	Showed Teklu Hadgu how to give user access to dropzone
	Troubleshoot reason some GoldSim jobs failed to return after 12 hours.

Page is intentionally blank

DISTRIBUTION

1	MS0736	Tito Bonano	08820
1	MS0747	Kevin McMahon	08842
1	MS0747	Robert J. MacKinnon	08844
1	MS0747	Gordon Appel	08842
1	MS0747	Teklu Hadgu	08842
1	MS0899	Technical Library	9536 (electronic copy)

Page is intentionally blank

